# Blind Separation of Nonlinear Mixtures by Variational Bayesian Learning 

Antti Honkela ${ }^{\text {a,* }}$ Harri Valpola ${ }^{\text {b }}$ Alexander Ilin ${ }^{\text {a }}$ Juha Karhunen ${ }^{\text {a }}$<br>${ }^{a}$ Adaptive Informatics Research Centre, Helsinki University of Technology P.O. Box 5400, FI-02015 TKK, Finland<br>${ }^{\mathrm{b}}$ Laboratory of Computational Engineering, Helsinki University of Technology P.O. Box 9203, FI-02015 TKK, Finland


#### Abstract

Blind separation of sources from nonlinear mixtures is a challenging and often illposed problem. We present three methods for solving this problem: an improved nonlinear factor analysis (NFA) method using multilayer perceptron (MLP) network to model the nonlinearity, a hierarchical NFA (HNFA) method suitable for larger problems and a post-nonlinear NFA (PNFA) method for more restricted post-nonlinear mixtures. The methods are based on variational Bayesian learning, which provides the needed regularisation and allows for easy handling of missing data. While the basic methods are incapable of recovering the correct rotation of the source space, they can discover the underlying nonlinear manifold and allow reconstruction of the original sources using standard linear independent component analysis (ICA) techniques.


Key words: Bayesian learning, blind source separation, nonlinear mixtures, post-nonlinear mixtures, variational Bayes PACS:

## 1 Introduction

Independent component analysis (ICA) [1], that is blind separation of independent sources from their linear mixtures, forms the basis for most source sep-

[^0]aration work today. Introduction of Bayesian methods based on the variational approximation for models capable of solving the ICA problem in 1999 independently by Attias [2] and Valpola ${ }^{1}$ [3] was an important step for Bayesian source separation. Later, slightly different methods based on the same principles have been proposed by a number of authors [4-7].

The Bayesian modelling approach is very useful in source separation, as in addition to providing tools for model order selection, it is open to all kinds of extensions simply by extending the model $[8,6,7]$. This was illustrated, for instance, by the nonlinear factor analysis (NFA) and nonlinear independent factor analysis (NIFA) models for blind separation of sources from nonlinear mixtures by Valpola and Honkela [9]. While the original NFA/NIFA model could be used to solve the difficult nonlinear separation problem in many cases, we have developed new variants of the model to better deal with certain specific instances.

The hierarchical nonlinear factor analysis (HNFA) method [10] was introduced to overcome the quadratic computational complexity with respect to the number of the sources in original NFA/NIFA, thus allowing solving larger problems. The special case of post-nonlinear mixtures was addressed by post-nonlinear factor analysis (PNFA) [11]. Other important general extensions include ability to handle missing data [12,13], improved initialisation strategies [14] as well as more accurate approximation of the nonlinearity in NFA/NIFA leading to increased stability [15]. In this paper we summarise these new results presented in scatterd conference papers, and present the improved methods in full detail together with new experimental comparisons.

The paper is organised as follows. In Sec. 2, the nonlinear blind source separation problem and some basic theory behind it are introduced, followed by a review of other existing methods in Sec. 3. This is followed by presentation of our models to solve the problem in Sec. 4. Learning the models is discussed in Sec. 5. Some experimental results on the methods are presented in Sec. 6 with discussion and conclusion in Secs. 7 and 8.

## 2 Nonlinear Blind Source Separation

Blind separation of sources from their nonlinear mixtures - known as nonlinear blind source separation (BSS) - is generally a difficult problem, both from a theoretical and a practical point of view $[1,16,17]$. The task is to extract the sources $\mathbf{s}(t)$ that have generated the observations $\mathbf{x}(t)$ through a nonlinear

[^1]mapping $\mathbf{f}(\cdot)$ :
\[

$$
\begin{equation*}
\mathbf{x}(t)=\mathbf{f}[\mathbf{s}(t)]+\mathbf{n}(t) \tag{1}
\end{equation*}
$$

\]

Here $\mathbf{x}(t)$ is the observed $N$-dimensional data (mixture) vector at time instant or index value $t, \mathbf{f}: \mathbb{R}^{M} \rightarrow \mathbb{R}^{N}$ is an unknown mixing function, $\mathbf{s}(t)$ is an $M$ vector whose elements are the $M$ unknown source signals $s_{i}(t), i=1,2, \ldots, M$, and $\mathbf{n}(t)$ is an $N$-dimensional vector containing the additive noise terms.

In the basic linear case

$$
\begin{equation*}
\mathbf{x}(t)=\mathbf{A} \mathbf{s}(t) \tag{2}
\end{equation*}
$$

where $\mathbf{A}$ is an $N \times M$ mixing matrix, the BSS problem can be solved using ICA [1] under certain conditions. Linear ICA methods utilise the strong but often realistic assumption that the source signals are statistically independent. Furthermore, all but possibly one of the source signals are assumed to be nonGaussian, there is no or only a little noise and the number of sources $M$ equals the number $N$ of mixtures [1]. Under these conditions, the sources in (2) can be estimated using ICA up to permutation, scaling, and sign.

Linear ICA and other BSS methods have been extended for the cases where the number of the sources is different from the number of mixtures, and for the noisy case. But the extension to the nonlinear mixture model (1) is difficult since both the nonlinear mapping and the underlying sources must be learned from the data in a blind manner, and the problem is highly ill-posed without suitable regularisation $[18,1,16]$. The independence assumption used in ICA is no longer sufficient for achieving separation. It can be shown that in the nonlinear case $x$ and $y$ can be mixed and still be statistically independent [18]. Additionally if $x$ and $y$ are two independent random variables, any of their functions $f(x)$ and $g(y)$ are also independent. These uniqueness issues have been considered in more detail in $[18,16]$.

A related problem is that it is often quite difficult to infer the number of sources and the structure of the mapping $\mathbf{f}(\cdot)$. From a practical point of view, efficiency and reliability of nonlinear BSS algorithms are critical issues. They have restricted the number of sources that can be separated in practice to be quite small in many instances. An important reason for this is that the computational load of many nonlinear BSS methods increases quite rapidly with the dimensionality of the problem, preventing in practice their application to high-dimensional data sets [1].

An important special case of the general nonlinear mixing model (1) consists of so-called post-nonlinear mixtures. There each mixture has the form

$$
\begin{equation*}
x_{i}(t)=f_{i}\left(\sum_{j=1}^{M} A_{i j} s_{j}(t)\right), \quad i=1, \ldots, N . \tag{3}
\end{equation*}
$$

Thus the sources $s_{j}(t), j=1, \ldots, M$, are first mixed linearly according to
the basic linear ICA/BSS model (2), but after that nonlinear functions $f_{i}$ are applied to them to get the final observations $x_{i}$. It can be shown $[19,16]$ that for the noiseless post-nonlinear mixtures, the indeterminacies are usually the same as for the basic linear instantaneous ICA model (2). That is, the sources can be separated or the independent components estimated up to the scaling, permutation, and sign indeterminacies assuming the post-nonlinearities $f_{i}$ are invertible and under weak conditions on the mixing matrix $\mathbf{A}$ and source distributions. The post-nonlinearity assumption is useful and reasonable in many signal processing applications, because it can be thought of as a model for a nonlinear sensor distortion. In more general situations, it is a restrictive and somewhat arbitrary constraint.

## 3 Existing methods

Existing nonlinear BSS methods have been reviewed in [16,17], and earlier in Chapter 17 of the book [1]. The paper [16] reviews especially uniqueness results on nonlinear ICA and BSS, separation methods for post-nonlinear mixtures, and our variational Bayesian estimation methods, referring to papers in which the detailed results have been presented. In this section, we will briefly review methods for BSS in nonlinear mixtures with i.i.d. samples.

The nonlinear blind source separation problem has been approached using auto-associative multilayer perceptron (MLP) networks, that is MLP networks trained with input-output pairs ( $\mathbf{x}, \mathbf{x}$ ) [20]. The number of nodes in a hidden layer is restricted to be smaller than the number of inputs and outputs, thus creating a bottleneck. The extracted nonlinear features can be retrieved from the values of the hidden nodes. With standard back-propagation this approach is very prone to overfitting and local minima, but more advanced learning methods such as flat minimum search can provide a method for nonlinear BSS [21]. The computational complexity from explicitly learning both mixing and demixing models can still be high.

The MISEP method [22] is a generalisation of the infomax method of linear ICA for nonlinear mixtures using an MLP network to model the nonlinearity. The method provides good separation results in several low-dimensional artificial examples as well as in a real nonlinear image mixture problem [23], but it has not been demonstrated for more than 4 mixtures of 4 sources. Nonlinear denoising source separation [24] can also be applied to separate low-dimensional nonlinear mixtures.

Kernel methods have recently become popular in producing nonlinear counterparts for many linear statistical methods based on second-order statistics. Kernel principal component analysis (KPCA) [25] is an extension of linear
principal component analysis (PCA) that operates on a high-dimensional feature space formed by a nonlinear transformation of the data space. The inner product features required for PCA can be evaluated efficiently in data space using the kernel trick. The method is computationally very efficient, but the results are highly dependent on the choice of the kernel.

The Bayesian kernel method of Gaussian Process Latent Variable Model (GPLVM) [26] offers a probabilistic alternative to KPCA. The presentation of the method seems to focus on use in visualisation and it has not been demonstrated for more than 3 dimensional latent spaces.

Post-nonlinear ICA has been approached, for instance, by direct minimisation of mutual information [19] as well as by heuristic methods. These are based on compensating for the invertible post-nonlinearities by first transforming the marginal densities of the data to be as Gaussian as possible and then applying standard ICA [27]. The downside of all of these methods is that they only work if the post-nonlinearities are invertible.

## 4 Models

In this section, the different models used in this work are presented. First, the multilayer perceptron (MLP) network based NFA model for general nonlinear mixtures from [9] is presented. Then, the hierarchical HNFA model based on the Bayes Blocks framework [33] and the PNFA model for post-nonlinear mixtures are introduced.

### 4.1 Multilayer Perceptron (MLP) Nonlinearity

Let us assume the observed data $\boldsymbol{X}$ follows the nonlinear mixing model of Eq. (1). The nonlinearity $\mathbf{f}$ is parametrised with a multilayer perceptron (MLP) network with single hidden layer with $H$ hidden nodes [20]. This allows for writing a generative model for a data vector $\mathbf{x}(t)$ as [9]

$$
\begin{equation*}
\mathbf{x}(t)=\mathbf{f}\left(\mathbf{s}(t), \boldsymbol{\theta}_{\mathbf{f}}\right)+\mathbf{n}(t)=\mathbf{B} \boldsymbol{\phi}(\mathbf{A} \mathbf{s}(t)+\mathbf{a})+\mathbf{b}+\mathbf{n}(t), \tag{4}
\end{equation*}
$$

where $\boldsymbol{\theta}_{\mathbf{f}}=(\mathbf{A}, \mathbf{B}, \mathbf{a}, \mathbf{b}), \mathbf{A} \in \mathbb{R}^{H \times M}$ and $\mathbf{B} \in \mathbb{R}^{N \times H}$ are the weight matrices and $\mathbf{a} \in \mathbb{R}^{H}$ and $\mathbf{b} \in \mathbb{R}^{N}$ are the bias vectors of the first and second layer of the MLP network, respectively. The weight matrices and bias vectors will hence be collectively called the weights of the MLP. The activation function $\phi$ is the standard hyperbolic tangent. It is applied component-wise to its argument vector. This approach is perfectly general because MLP networks are universal
function approximators that can model any nonlinear mapping to arbitrary accuracy [28].

The noise $\mathbf{n}(t)$ and all the weight matrices and bias vectors of the MLP are a priori assumed to be Gaussian and independent of each other. The noise is assumed to have a general diagonal covariance and a hierarchical log-normal variance prior of the form

$$
\begin{align*}
p\left(\mathbf{n}(t) \mid \mathbf{v}_{n}\right) & =N\left(\mathbf{n}(t) ; \mathbf{0}, \operatorname{diag}\left(\exp \left(2 \mathbf{v}_{n}\right)\right)\right)  \tag{5}\\
p\left(v_{n_{i}} \mid m_{v_{n}}, v_{v_{n}}\right) & =N\left(v_{n_{i}} ; m_{v_{n}}, \exp \left(2 v_{v_{n}}\right)\right), \tag{6}
\end{align*}
$$

where $\mathbf{v}_{n}=\left(v_{n_{1}}, \ldots, v_{n_{N}}\right)$ and $\operatorname{diag}(\mathbf{v})$ is a diagonal matrix with the elements of vector $\mathbf{v}$ on the main diagonal. The log-normal parametrisation is used because it is easy to define a hierarchical prior. As the variance model is not that important in this application, a conjugate inverse Gamma variance prior could be used here as well. Restriction to an isotropic noise model with the noise covariance of the form $\lambda \mathbf{I}$ is also possible. The noise model and the generative model (4) imply the likelihood

$$
\begin{equation*}
p\left(\mathbf{x}(t) \mid \boldsymbol{\theta}_{\mathbf{f}}, \mathbf{s}(t), \mathbf{v}_{n}\right)=N\left(\mathbf{x}(t) ; \mathbf{f}\left(\mathbf{s}(t), \boldsymbol{\theta}_{\mathbf{f}}\right), \operatorname{diag}\left(\exp \left(2 \mathbf{v}_{n}\right)\right)\right) . \tag{7}
\end{equation*}
$$

The hierarchical priors of the MLP weights are similar to those of the noise except that the prior of $\mathbf{A}$ is fixed to unit variance to resolve the scaling ambiguity between $\mathbf{A}$ and $\mathbf{s}$, and the different columns of $\mathbf{B}$ have their own priors:

$$
\begin{align*}
p\left(A_{i j}\right) & =N\left(A_{i j} ; 0,1\right)  \tag{8}\\
p\left(B_{i j} \mid v_{B_{j}}\right) & =N\left(B_{i j} ; 0, \exp \left(2 v_{B_{j}}\right)\right)  \tag{9}\\
p\left(a_{i} \mid m_{a}, v_{a}\right) & =N\left(a_{i} ; m_{a}, \exp \left(2 v_{a}\right)\right)  \tag{10}\\
p\left(b_{i} \mid m_{b}, v_{b}\right) & =N\left(b_{i} ; m_{b}, \exp \left(2 v_{b}\right)\right)  \tag{11}\\
p\left(v_{B_{j}} \mid m_{v_{B}}, v_{v_{B}}\right) & =N\left(v_{B_{j}} ; m_{v_{B}}, \exp \left(2 v_{v_{B}}\right)\right) . \tag{12}
\end{align*}
$$

The highest-level hyperparameters $m_{v_{n}}, v_{v_{n}}, m_{a}, v_{a}, m_{b}, v_{b}, m_{v_{B}}$ and $v_{v_{B}}$ have vague Gaussian priors $N\left(0,100^{2}\right)^{2}$.

To complete the definition of the model, the prior of the sources $\mathbf{s}$ must still be determined. In this case there are several possibilities, leading to models with different benefits and drawbacks. The differences between the alternatives are discussed in more detail in Section 4.4 below.

The simplest source model is the Gaussian model used in nonlinear factor
$\overline{2}$ The data is usually preprocessed by scaling to approximately unit variance to ensure proper scaling of the weights so that the priors really are vague.
analysis (NFA). This is achieved by

$$
\begin{align*}
p\left(\mathbf{s}(t) \mid \mathbf{v}_{s}\right) & =N\left(\mathbf{s}(t) ; \mathbf{0}, \operatorname{diag}\left(\exp \left(2 \mathbf{v}_{s}\right)\right)\right)  \tag{13}\\
p\left(v_{s_{i}} \mid m_{v_{s}}, v_{v_{s}}\right) & =N\left(v_{s_{i}} ; m_{v_{s}}, \exp \left(2 v_{v_{s}}\right)\right), \tag{14}
\end{align*}
$$

where the hyperparameters $m_{v_{s}}$ and $v_{v_{s}}$ again have noninformative prior $N\left(0,100^{2}\right)$. With these definitions, the sets of observations $\boldsymbol{X}=\{\mathbf{x}(t) \mid t\}$ and sources $\boldsymbol{S}=$ $\{\mathbf{s}(t) \mid t\}$ are defined as usual. The parameter vector $\boldsymbol{\theta}$ contains everything described above, that is $\boldsymbol{\theta}=\left(\boldsymbol{\theta}_{\mathbf{f}}, \mathbf{v}_{n}, m_{v_{n}}, v_{v_{n}},\left(v_{B_{j}}\right), m_{a}, v_{a}, m_{b}, v_{b}, m_{v_{B}}, v_{v_{B}}, \mathbf{v}_{s}, m_{v_{s}}, v_{v_{s}}\right)$ including all $j$ in $v_{B_{j}}$.

The Gaussian model is computationally simple, but it suffers from the same rotation indeterminacy as the linear factor analysis (FA) model. This can be corrected in the same way linear independent factor analysis [2] extends linear FA: by using a mixture of Gaussians as the source prior. Introducing a new latent variable $M_{i}(t)$ to denote the active mixture component for source $i$ at sample $t$ leads to a nonlinear independent factor analysis (NIFA) model with a prior of the form

$$
\begin{align*}
p\left(s_{i}(t) \mid M_{i}(t)=l, m_{s i l}, v_{s i l}\right) & =N\left(s_{i}(t) ; m_{\text {sil }}, \exp \left(2 v_{s i l}\right)\right)  \tag{15}\\
p\left(m_{s i l} \mid v_{m_{s}}\right) & =N\left(m_{s i l} ; 0, \exp \left(2 v_{m_{s}}\right)\right)  \tag{16}\\
p\left(v_{s i l} \mid m_{v_{s}}, v_{v_{s}}\right) & =N\left(v_{s i l} ; m_{v_{s}}, \exp \left(2 v_{v_{s}}\right)\right) . \tag{17}
\end{align*}
$$

The mixing proportions have a logistic normal prior given by the softmax function

$$
\begin{align*}
p\left(M_{i}(t)=l \mid c_{i \cdot}\right) & =\exp \left(c_{i l}\right) / \sum_{l^{\prime}} \exp \left(c_{i l^{\prime}}\right)  \tag{18}\\
p\left(c_{i l} \mid v_{c}\right) & =N\left(c_{i l} ; 0, \exp \left(2 v_{c}\right)\right) . \tag{19}
\end{align*}
$$

All the highest level parameters $v_{m_{s}}, m_{v_{s}}, v_{v_{s}}$ and $v_{c}$ again have a noninformative prior $N\left(0,100^{2}\right)$. The parameters $\boldsymbol{\theta}$ can be defined similarly as above.

In the following, the simpler NFA model is used instead of the more complex NIFA model. More detailed justification for this is presented in Sec. 4.4. Most of the discussion generalises fairly easily to NIFA as well, as is shown in [9,29].

### 4.2 Hierarchical Nonlinearity

The computational complexity of learning the NFA model is quadratic with respect to the number of the sources. This arises from the need to evaluate the Jacobian matrix of the outputs of the MLP network with respect to the inputs for each sample. The Jacobian is needed to reliably approximate the nonlinearity [15].

One way of defining a more efficiently learnable new model is to introduce additional latent variables to the hidden nodes of the MLP-like network as is done in the hierarchical nonlinear factor analysis (HNFA) model [10]. The HNFA model is defined by the equations

$$
\begin{align*}
\mathbf{s}(t) & \sim N(\mathbf{0}, \mathbf{I})  \tag{20}\\
\mathbf{h}(t) & \sim N\left(\mathbf{A s}(t)+\mathbf{a}, \operatorname{diag}\left(\exp \left(-\mathbf{v}_{h}\right)\right)\right)  \tag{21}\\
\mathbf{x}(t) & \sim N\left(\mathbf{C s}(t)+\mathbf{B} \boldsymbol{\phi}(\mathbf{h}(t))+\mathbf{b}, \operatorname{diag}\left(\exp \left(-\mathbf{v}_{n}\right)\right)\right) \tag{22}
\end{align*}
$$

where $\mathbf{h}(t)$ are the latent variables modelling the values of the hidden nodes, the vectors $\mathbf{v}_{h}$ and $\mathbf{v}_{n}$ parametrise the variances of the noise or innovation of the hidden nodes and the observations, respectively, $\boldsymbol{\phi}$ is a vector of activation functions and $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{a}, \mathbf{b}$ are the weights of the mapping. The activation function is chosen to be $\phi(y)=\exp \left(-y^{2}\right)$, for computational simplicity. The structure of the model is illustrated in Fig. 1.


Fig. 1. The HNFA model is illustrated. Square nodes correspond to weight matrices and round nodes to the variables with shaded nodes being observed and unshaded latent.

When assigning hierarchical priors to the parameters of the mapping, we have taken into account the role of different parameters in the model. The weights between different layers, that is $\mathbf{A}, \mathbf{B}$ and $\mathbf{C}$, can be expected to scale according to the scales of inputs and outputs. Since all $s_{i}(t)$ are assumed normally distributed, there should be no significant differences in the scales of the weights leaving different nodes but the hidden nodes and observations can
have different scales. This is reflected in the following parametrisation:

$$
\begin{align*}
p\left(A_{i j} \mid v_{I, h_{i}}\right) & =N\left(A_{i j} ; 0, \exp \left(-v_{I, h_{i}}\right)\right)  \tag{23}\\
p\left(B_{i j} \mid v_{I, x_{i}}, v_{O, h_{j}}\right) & =N\left(B_{i j} ; 0, \exp \left(-v_{I, x_{i}}-v_{O, h_{j}}\right)\right)  \tag{24}\\
p\left(C_{i j} \mid v_{I, x_{i}}\right) & =N\left(C_{i j} ; 0, \exp \left(-v_{I, x_{i}}\right)\right)  \tag{25}\\
p\left(a_{i}\right) & =N\left(a_{i} ; 0,1\right)  \tag{26}\\
p\left(b_{i}\right) & =N\left(b_{i} ; 0,1\right) . \tag{27}
\end{align*}
$$

The parameters $v_{O, h_{j}}$ reflect the different scales of weights leaving the hidden nodes and $v_{I, h_{i}}$ and $v_{I, x_{i}}$ correspond to the different scales of weights entering the hidden nodes and observations, respectively.

Finally, the parameters which model variances on logarithmic scale, that is $\mathbf{v}_{h}$, $\mathbf{v}_{n}, v_{I, h_{i}}, v_{I, x_{i}}$ and $v_{O, h_{j}}$, each have hyperparameters for mean and log-variance. Their priors are $N(0, \exp (7))$ and $N(0, \exp (4))$, respectively.

MLP networks with sigmoid activation function are known to be universal approximators [28]. Similarly, the set of deterministic HNFA mappings without the linear shortcut

$$
\begin{equation*}
f(\mathbf{s})=\mathbf{B} \boldsymbol{\phi}(\mathbf{A s}(t)+\mathbf{a})+\mathbf{b} \tag{28}
\end{equation*}
$$

where $\phi(y)=\exp \left(-y^{2}\right)$, has the universal approximation property exactly analogous to that of MLP networks [30].

The HNFA model can be implemented efficiently using the building-blocks framework [31-33]. As the additional latent variables $\mathbf{h}(t)$ are independent by definition of the factorial posterior approximation, there is no need to evaluate any nontrivial Jacobians. The disadvantage of HNFA is that the approximation of the posterior density is farther away from the true posterior density. This may occasionally lead to inferior performance [34]. The linear shortcut is included in the model to partially help this, because representing nearly linear mappings would otherwise be significantly more difficult than with the MLP model.

### 4.3 Post-nonlinear Model

As discussed in Section 2, post-nonlinear (PNL) mixtures are restricted nonlinear mixtures of the form

$$
\begin{equation*}
x_{i}(t)=f_{i}\left(\sum_{j=1}^{M} A_{i j} s_{j}(t)\right)+n_{i}(t), \quad i=1, \ldots, N \tag{29}
\end{equation*}
$$

where the functions $f_{i}: \mathbb{R} \rightarrow \mathbb{R}$ are called the post-nonlinearities and $n_{i}$ is additive Gaussian noise. PNL mixtures are a theoretically important special
case of nonlinear ICA, because they can be proven to be separable in the noiseless case if the post-nonlinearities are invertible and the mixing matrix satisfies certain regularity conditions [19].

A variational Bayesian approach to post-nonlinear mixtures was first presented in [11]. The post-nonlinear factor analysis (PNFA) model is based on the generative model

$$
\begin{equation*}
x_{i}(t)=f_{i}\left(\sum_{j=1}^{M} A_{i j} s_{j}(t) ; \boldsymbol{\theta}_{f_{i}}\right)+n_{i}(t), \tag{30}
\end{equation*}
$$

where the post-nonlinearities $f_{i}$ are modelled with MLP networks with weights $\boldsymbol{\theta}_{f_{i}}$ as

$$
\begin{equation*}
f_{i}\left(y, \boldsymbol{\theta}_{f_{i}}\right)=\mathbf{D}_{i} \boldsymbol{\phi}\left(\mathbf{C}_{i} y+\mathbf{c}_{i}\right)+d_{i}, \tag{31}
\end{equation*}
$$

where $\boldsymbol{\theta}_{f_{i}}=\left(\mathbf{C}_{i}, \mathbf{c}_{i}, \mathbf{D}_{i}, d_{i}\right)$ includes the column weight and bias vectors $\mathbf{C}_{i}, \mathbf{c}_{i}$, row weight vector $\mathbf{D}_{i}$ and bias scalar $d_{i}$. This approach also allows for noninvertible post-nonlinearities.

Most of the weights have priors governed by higher level hyperparameters

$$
\begin{align*}
p\left(A_{i j}\right) & =N\left(A_{i j} ; 0,1\right)  \tag{32}\\
p\left(C_{i j} \mid v_{C_{i}}\right) & =N\left(C_{i j} ; 0, \exp \left(2 v_{C_{i}}\right)\right)  \tag{33}\\
p\left(c_{i j} \mid m_{c_{i}}, v_{c_{i}}\right) & =N\left(c_{i j} ; m_{c_{i}}, \exp \left(2 v_{c_{i}}\right)\right)  \tag{34}\\
p\left(D_{i j} \mid v_{D_{i}}\right) & =N\left(D_{i j} ; 0, \exp \left(2 v_{D_{i}}\right)\right)  \tag{35}\\
p\left(d_{i} \mid m_{d}, v_{d}\right) & =N\left(d_{i} ; m_{d}, \exp \left(2 v_{d}\right)\right), \tag{36}
\end{align*}
$$

with hierarchical priors for $v_{C_{i}}, m_{c_{i}}, v_{c_{i}}, v_{D_{i}}$ similar to those used in the NFA model in Eq. (12). The highest level parameters again have vague priors $N\left(0,100^{2}\right)$.

The source model used in PNFA is Gaussian. It could in principle be rather easily replaced with mixtures-of-Gaussians, but simple post-processing of the extracted sources with linear ICA is often sufficient, as discussed below in Sec. 4.4.

### 4.4 On different source models

Two different source models were presented above: a Gaussian nonlinear factor analysis (NFA) model in Eq. (13) and a mixture-of-Gaussians based nonlinear independent factor analysis (NIFA) model in Eq. (15). From purely theoretical perspective, the NIFA model is preferable as its non-Gaussian model is able to resolve the rotation indeterminacy inherent to the Gaussian model.

In practice, things are not quite that simple. The NIFA model with its fully factorial posterior approximation seems to suffer from similar problems of not being able to extract the true independent sources as linear algorithms based on similar principles [34].

Using the more complicated NIFA model slows down the computation and requires careful initialisation of the mixture components. As this seems to cause more trouble than offer benefits, we have usually chosen not to do it and only use plain NFA. In order to achieve BSS, standard linear ICA can be applied as post-processing to the extracted sources. This approach resembles the one used in [35] for linear separation, where linear factor analysis is used as pre-processing for ICA instead of PCA. Use of factor analysis pre-processing has been shown to lead to good results in case of noisy linear mixtures. The nonlinear case is, however, more difficult as the model may try to nonlinearly transform the sources to be more Gaussian than they should be. Some evidence of this behaviour can be seen in the experiments in Sec. 6.2. Even with a Gaussian source model the sources are, however, not fully Gaussian and using linear ICA to determine the rotation is thus possible.

## 5 Learning

In this section, methods for learning the models presented in Sec. 4 are presented. The learning method for NFA is again heavily based on earlier work [9] but uses a new method of linearising the MLP [15] that leads to significantly more stable and reliable algorithm.

All the models presented in Sec. 4 are learned using variational Bayesian (VB) learning [36]. Given the data $\boldsymbol{X}$ and a model with parameters $\boldsymbol{\theta}$ and latent variables or sources $\boldsymbol{S}$, the method is based on approximating the posterior distribution $p(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{X}, \mathcal{H})$ with a factorial approximation $q(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{\xi})=$ $q\left(\boldsymbol{S} \mid \boldsymbol{\xi}_{\boldsymbol{S}}\right) q\left(\boldsymbol{\theta} \mid \boldsymbol{\xi}_{\boldsymbol{\theta}}\right)$, where $\boldsymbol{\xi}$ are the variational parameters and $\mathcal{H}$ represents the background assumptions of a particular model. The approximation is fitted by minimising the Kullback-Leibler divergence between $q$ and $p$, or equivalently maximising a lower bound on marginal log-likelihood of the data. By changing the sign, this can be formulated as minimisation of the cost

$$
\begin{align*}
\mathcal{C} & =\left\langle\log \frac{q(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{\xi})}{p(\boldsymbol{S}, \boldsymbol{\theta}, \boldsymbol{X} \mid \mathcal{H})}\right\rangle  \tag{37}\\
& =D_{K L}(q(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{\xi})| | p(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{X}, \mathcal{H}))-\log p(\boldsymbol{X} \mid \mathcal{H}) \geq-\log p(\boldsymbol{X} \mid \mathcal{H})
\end{align*}
$$

where $\langle\cdot\rangle$ denotes the expectation over $q$ and $D_{K L}(q \| p)$ is the Kullback-Leibler divergence between $q$ and $p$.

Most applications of variational Bayesian learning [37] involve only models in the conjugate exponential family. The models described in Sec. 4 are clearly not in the exponential family, so the learning methods are somewhat different and require numerical minimisation of the cost or even only its approximation.

In almost all the models, we use the naïve mean field approximation, that is a fully factorial Gaussian distribution

$$
\begin{equation*}
q(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{\xi})=q\left(\boldsymbol{S} \mid \boldsymbol{\xi}_{S}\right) q\left(\boldsymbol{\theta} \mid \boldsymbol{\xi}_{\boldsymbol{\theta}}\right)=\prod_{i, t} q\left(s_{i}(t) \mid \boldsymbol{\xi}_{s_{i}(t)}\right) \prod_{j} q\left(\theta_{j} \mid \boldsymbol{\xi}_{\theta_{j}}\right) \tag{38}
\end{equation*}
$$

The individual factors are parametrised with variational parameters corresponding to the posterior mean and variance of the variable as

$$
\begin{align*}
q\left(s_{i}(t) \mid \boldsymbol{\xi}_{s_{i}(t)}\right) & =N\left(s_{i}(t) ; \bar{s}_{i}(t), \widetilde{s}_{i}(t)\right)  \tag{39}\\
q\left(\theta_{j} \mid \boldsymbol{\xi}_{\theta_{j}}\right) & =N\left(\theta_{j} ; \bar{\theta}_{j}, \widetilde{\theta}_{j}\right) . \tag{40}
\end{align*}
$$

For parameters modelling the mean of a Gaussian, the Gaussian distribution is a conjugate prior and the optimal free-form approximation is of this form, assuming the factorisation. For parameters modelling the log-variance of a Gaussian, the Gaussian prior is not conjugate and the posterior approximation is only an approximation.

The definition of the model and the approximating distribution allow evaluating the cost (37) as a function of the variational parameters $\boldsymbol{\xi}$. The cost can be split into two parts as the negative entropy of the approximation and an expected energy

$$
\begin{equation*}
\mathcal{C}=\mathcal{C}_{q}+\mathcal{C}_{p}=\langle\log q(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{\xi})\rangle+\langle-\log p(\boldsymbol{X}, \boldsymbol{S}, \boldsymbol{\theta} \mid \mathcal{H})\rangle \tag{41}
\end{equation*}
$$

Using the factorisation of $q$, the negative entropy term $\mathcal{C}_{q}$ splits into

$$
\begin{align*}
\mathcal{C}_{q} & =\langle\log q(\boldsymbol{S}, \boldsymbol{\theta} \mid \boldsymbol{\xi})\rangle=\left\langle\log \prod_{i, t} q\left(s_{i}(t) \mid \boldsymbol{\xi}_{s_{i}(t)}\right) \prod_{j} q\left(\theta_{j} \mid \boldsymbol{\xi}_{\theta_{j}}\right)\right\rangle  \tag{42}\\
& =\sum_{i, t}\left\langle\log q\left(s_{i}(t) \mid \boldsymbol{\xi}_{s_{i}(t)}\right)\right\rangle+\sum_{j}\left\langle\log q\left(\theta_{j} \mid \boldsymbol{\xi}_{\theta_{j}}\right)\right\rangle .
\end{align*}
$$

The remaining terms are negative entropies of univariate Gaussians having values depending only on the variance

$$
\begin{equation*}
\left\langle\log q\left(\theta_{j} \mid \boldsymbol{\xi}_{\theta_{j}}\right)\right\rangle=-\frac{1}{2}-\frac{1}{2} \log \left(2 \pi \widetilde{\theta}_{j}\right) . \tag{43}
\end{equation*}
$$

The expected energy term $\mathcal{C}_{p}$ is more model specific and is presented separately for the different models.

### 5.1 MLP nonlinearity

While basic NFA uses a fully factorial posterior approximation, this is not the case for NIFA. Instead, the dependences between $M_{i}(t)$ and $s_{i}(t)$ are modelled so that the approximation for them is of the form

$$
\begin{equation*}
q\left(M_{i}(t), s_{i}(t)\right)=q\left(s_{i}(t) \mid M_{i}(t)\right) q\left(M_{i}(t)\right) \tag{44}
\end{equation*}
$$

which yields a Gaussian mixture approximation for $s_{i}(t)$

$$
\begin{equation*}
q\left(s_{i}(t)\right)=\sum_{l} q\left(s_{i}(t) \mid M_{i}(t)=l\right) q\left(M_{i}(t)=l\right) \tag{45}
\end{equation*}
$$

Otherwise the approximation is similar to NFA.

### 5.1.1 Evaluating the cost

Evaluating the expected energy term $\mathcal{C}_{p}$ is slightly more difficult than term $\mathcal{C}_{q}$. Using the definition of the model, it can for NFA and PNFA models be factored into terms

$$
\begin{align*}
\mathcal{C}_{p} & =\langle-\log p(\boldsymbol{X} \mid \boldsymbol{S}, \boldsymbol{\theta}, \mathcal{H})\rangle+\langle-\log p(\boldsymbol{S} \mid \boldsymbol{\theta}, \mathcal{H})\rangle+\langle-\log p(\boldsymbol{\theta} \mid \mathcal{H})\rangle \\
& =\sum_{i, t}\left\langle-\log p\left(x_{i}(t) \mid \mathbf{s}(t), \boldsymbol{\theta}, \mathcal{H}\right)\right\rangle+\sum_{i, t}\left\langle-\log p\left(s_{i}(t) \mid \boldsymbol{\theta}, \mathcal{H}\right)\right\rangle+\langle-\log p(\boldsymbol{\theta} \mid \mathcal{H})\rangle \tag{46}
\end{align*}
$$

The terms in the second and third summand are expectations of negative logarithm of Gaussian pdf over Gaussian mean and log-variance parameters. These can be evaluated for example parameter $\theta \sim N(m, \exp (2 v))$ through integrals of the form [38]

$$
\begin{align*}
\langle-\log p(\theta \mid m, v, \mathcal{H})\rangle & =\iiint-\log N(\theta ; m, \exp (2 v)) q(\theta) q(m) q(v) d \theta d m d v \\
& =\frac{1}{2} \log (2 \pi)+\bar{v}+\frac{1}{2}\left[(\bar{\theta}-\bar{m})^{2}+\widetilde{\theta}+\widetilde{m}\right] \exp (2 \widetilde{v}-2 \bar{v}) \tag{47}
\end{align*}
$$

The terms of the first summand of Eq. (46),

$$
\begin{align*}
\left\langle-\log p\left(x_{i}(t) \mid \mathbf{s}(t), \boldsymbol{\theta}, \mathcal{H}\right)\right\rangle & =\left\langle-\log N\left(x_{i}(t) ; f_{i}\left(\mathbf{s}(t), \boldsymbol{\theta}_{\mathbf{f}}\right), \exp \left(v_{n_{i}}\right)\right)\right\rangle \\
=\frac{1}{2} \log (2 \pi)+\bar{v}_{n_{i}} & +\frac{1}{2}\left[\left(x_{i}(t)-\bar{f}_{i}(t)\right)^{2}+\widetilde{f}_{i}(t)\right] \exp \left(2 \widetilde{v}_{n_{i}}-2 \bar{v}_{n_{i}}\right), \tag{48}
\end{align*}
$$

are more difficult as they depend on the mean $\bar{f}_{i}(t)$ and variance $\tilde{f}_{i}(t)$ of the outputs of the MLP network. Techniques for approximating these are presented in Appendix A.

### 5.1.2 Update algorithm

The posterior approximations of the parameters of the hierarchical model, that is those of the type $m_{\theta}$ and $v_{\theta}$, can be updated using a standard VB EM algorithm $[38,9]$. The update rules for the sources $\boldsymbol{S}$ as well as the weights $\boldsymbol{\theta}_{\mathbf{f}}$ of the MLP network are more difficult and they are therefore presented in more detail ${ }^{3}$.

The variational parameters corresponding to the means of $\boldsymbol{S}$ and $\boldsymbol{\theta}_{\mathrm{f}}$ are updated with a conjugate gradient algorithm. The required derivatives can be computed from the expression of the cost function. In case of inputs and weights of the MLP, this leads to similar computation of partial derivatives as in back-propagation. The parameters corresponding to the variances are updated using a fixed point algorithm that can be derived by differentiating the split cost (41) with respect to $\widetilde{\theta}_{j}$ and using the evaluated result (43). This yields

$$
\begin{equation*}
\frac{\partial \mathcal{C}}{\partial \widetilde{\theta}_{j}}=\frac{\partial \mathcal{C}_{q}}{\partial \tilde{\theta}_{j}}+\frac{\partial \mathcal{C}_{p}}{\partial \widetilde{\theta}_{j}}=-\frac{1}{2 \widetilde{\theta}_{j}}+\frac{\partial \mathcal{C}_{p}}{\partial \tilde{\theta}_{j}} \tag{49}
\end{equation*}
$$

Setting this to zero leads to a fixed point update rule for the variances of the sources and MLP network weights

$$
\begin{equation*}
\widetilde{\theta}_{j}=\frac{1}{2}\left(\frac{\partial \mathcal{C}_{p}}{\partial \tilde{\theta}_{j}}\right)^{-1} . \tag{50}
\end{equation*}
$$

Blindly applying this rule may in some cases lead to instability. This can be corrected by some form of dampening, such as halving the step length on logarithmic scale until the update does not increase the value of the cost function. The variance must also not be set to a negative value, even if the derivative is negative.

### 5.1.3 Initialisation

The MLP network and the gradient-based learning algorithms are notoriously prone to local minima $[40,20]$. In order to achieve good results, the methods require a reasonable initialisation.

The NFA method is typically initialised by setting the means of the sources to values given by suitable number of principal components of the data. The means of the MLP weights are initialised to random values while all the variances are initialised to small constant values. After this, only the MLP weights are updated during the first 20 iterations of the update algorithm so that the

[^2]model can learn a mapping from the PCA sources to the observations. The hyperparameters of the model are only updated after the first 100 iterations.

The PCA initialisation is easy to compute and sufficient for many purposes, but its linearity may sometimes lead to suboptimal results. To resolve this, kernel PCA (KPCA) was used in the initialisation in [14]. The KPCA initialisations are sensitive to the choice of the kernel and its parameters, but with suitable choices they may help NFA yield significantly better results while using less time. The kernel can also be selected with the variational Bayesian criterion by running the learning algorithm for a few iterations with different initialisations and comparing the cost function values.

### 5.2 Hierarchical nonlinearity

The HNFA model is defined by a graphical model of Gaussian variables combined using addition, multiplication and nonlinearities directly following the variables. Such models can be implemented using the building-blocks framework introduced in [33]. The learning scheme is designed to minimise the cost function

$$
\begin{equation*}
\mathcal{C}=\left\langle\log \frac{q(\boldsymbol{S}, \boldsymbol{H}, \boldsymbol{\theta} \mid \boldsymbol{\xi})}{p(\boldsymbol{S}, \boldsymbol{H}, \boldsymbol{\theta}, \boldsymbol{X} \mid \mathcal{H})}\right\rangle, \tag{51}
\end{equation*}
$$

where $\boldsymbol{H}=\{\mathbf{h}(t) \mid t\}$ is the set of all hidden nodes. The cost function is otherwise the same as Eq. (37), except that the hidden nodes $\boldsymbol{H}$ have been included as unknown variables.

The basic operation during learning is an iteration where all the terms $q\left(\theta_{i} \mid \boldsymbol{\xi}_{\theta_{i}}\right)$, $q\left(s_{i}(t) \mid \boldsymbol{\xi}_{s_{i}(t)}\right)$ and $q\left(h_{i}(t) \mid \boldsymbol{\xi}_{h_{i}(t)}\right)$ of the approximation

$$
\begin{align*}
q(\boldsymbol{S}, \boldsymbol{\theta}, \boldsymbol{H} \mid \boldsymbol{\xi}) & =q\left(\boldsymbol{S} \mid \boldsymbol{\xi}_{\boldsymbol{S}}\right) q\left(\boldsymbol{\theta} \mid \boldsymbol{\xi}_{\boldsymbol{\theta}}\right) q\left(\boldsymbol{H} \mid \boldsymbol{\xi}_{\boldsymbol{H}}\right) \\
& =\prod_{i, t} q\left(s_{i}(t) \mid \boldsymbol{\xi}_{s_{i}(t)}\right) \prod_{j} q\left(\theta_{j} \mid \boldsymbol{\xi}_{\theta_{j}}\right) \prod_{i, t} q\left(h_{i}(t) \mid \boldsymbol{\xi}_{h_{i}(t)}\right) \tag{52}
\end{align*}
$$

are updated one at a time by minimising (51) with respect to the corresponding variational parameters $\boldsymbol{\xi}$. The approximation is again otherwise the same as Eq. (38), except for the inclusion of the hidden nodes $\boldsymbol{H}$. Details of the minimisation of the $q$-terms are presented in [33, Appendices A and C] ${ }^{4}$.

In addition to the basic updates, several other operations are performed in order to help avoid local optima and speed up learning:

- addition of hidden nodes;
$\overline{4}$ The Bayes Blocks toolbox [41] and the code for the HNFA method are available at http://www.cis.hut.fi/projects/bayes/software/.
- addition of weights;
- pruning of weights; and
- line search.

Line search has been explained in [42]. The idea is to monitor the individual updates during one iteration and then perform a line search simultaneously for all terms of the approximation. We applied the line search after every tenth iteration.

The addition and pruning operations aim at optimising the model structure. Incremental addition of hidden nodes to the HNFA model is necessary because the hidden nodes are "expensive" and thus easily pruned out by making their outgoing weights approach to zero. The cost function (37) relates to the marginal likelihood $p(\boldsymbol{X} \mid \mathcal{H})$ which can be used to find the most likely model structure.

In general, addition takes place randomly and pruning is based on estimating whether the cost function can be decreased by removing a weight. The motivation for this is that variational Bayesian learning can effectively prune out parts of the model which are not needed. The weights in the matrix $\mathbf{B}$ corresponding to one hidden node can for instance approach zero. The cost function can usually be decreased by removing such weights. If all outgoing weights of a hidden node have been removed, the hidden node becomes useless and can be removed from the model. Variational Bayesian learning cannot, however, actively make room for a part of the model which may be added in the future. It usually takes some time for the rest of the model to accommodate to additions.

### 5.2.1 Evidence node

In order to avoid local minima in variational Bayesian learning, it is necessary to initialise some variables and keep them fixed for a while until other parts of the model have accommodated appropriately. In the building-blocks framework, we have adopted the virtual-evidence approach [43]. This is implemented by attaching evidence nodes [44], as we call them, to a variable $\theta_{i}$, whose value we want to set. The node provides a term for the likelihood $p$ (children $\left.\mid \theta_{i}\right)$. When $q\left(\theta_{i} \mid \boldsymbol{\xi}_{\theta_{i}}\right)$ is updated, the distribution of $\theta_{i}$ will be close to the value set by evidence node if the likelihood term has a narrow peak but $\theta_{i}$ can accommodate to other parts of the model if the likelihood term is wide. Since all the prior distributions in the present model are Gaussian, it is convenient to use the same functional form for the evidence nodes:

$$
\begin{equation*}
-\log p\left(\text { virtual child } \mid \theta_{i} ; \alpha_{i}\left(\tau_{i}\right), \mu_{i}, \sigma_{i}^{2}\right)=\alpha_{i}\left(\tau_{i}\right) \frac{\left(\theta_{i}-\mu_{i}\right)^{2}}{2 \sigma_{i}^{2}} \tag{53}
\end{equation*}
$$

where $\mu_{i}$ and $\sigma_{i}^{2}$ are the mean and variance induced by the evidence node. The term $\alpha\left(\tau_{i}\right)$ controls the decay of the virtual evidence. It decreases linearly from $\alpha(0)=1$ to $\alpha(T)=0$ after which the evidence node is removed. The persistence of the initialisation can thus be controlled by the life-span $T$ (iterations) of the evidence node.

### 5.2.2 Phases of learning

The model is built in stages. First, only the linear mapping $\mathbf{C s}(t)$ is used, so that there are no hidden nodes. The sources $\mathbf{s}(t)$ are initialised by principal component analysis (PCA) using evidence nodes with a life-span of 40 iterations and variance $\sigma^{2}=10^{-3}$ in order to estimate a reasonable $\mathbf{C}$. The linear model is learned for 100 iterations.

After that, 50 randomly initialised hidden nodes are added to the model and estimation of the model structure begins. That is, weights are added and pruned and hidden nodes are added every now and then. Every time new hidden nodes are added, five of them are selected from a pool of 1,000 random candidates. After each addition of hidden nodes, there is a period of 30 iterations during which no pruning is applied. This gives the new hidden nodes enough time to fit themselves into the model.

Hidden nodes are added a limited number of times. After that, learning continues with pruning and random additions of weights. The number of weights to be added decreases with time. Finally, only line searches are applied for the last 1,000 iterations. The total number of iterations in the HNFA simulations is 10,000 unless otherwise stated.

### 5.2.3 Addition of hidden nodes

The hidden nodes are latent variables which can independently represent some aspects of the observations. Due to our model structure, this usually corresponds to a local minimum of the cost function. It is better that the sources $\mathbf{s}(t)$ represent the data since they can share their information with all hidden nodes. Local minima can be avoided by evidence nodes which keep the variance of the Gaussian noise, innovation, associated to each newly added hidden node $h_{i}(t)$ through the corresponding parameter $v_{h_{i}}$ low. Once the sources take the responsibility for the representation, the variances of hidden nodes no longer grow significantly. In our experiments, the life-span of these evidence nodes was 500 iterations and their mean and variance were set to 9 and unity, respectively. Since the parameter controls the logarithm of the inverse variance, value of 9 corresponds to standard deviation of approximately $1 / 90$.

Fig. 2 shows a clear correlation of the quality of the results for the experimental
setting described in Sec. 6.1, and a measure of the amount of independent innovation in the hidden nodes, the latter of which can be influenced by introducing the evidence nodes. More detailed discussion is presented in [45].


Fig. 2. The attained signal-to-noise ratio of the sources separated using HNFA with and without evidence nodes. (From [45].)

In addition to restricting the innovations, the incoming weights $\mathbf{A}$ of the hidden nodes are initialised to random values by evidence nodes with variance $\sigma^{2}=10^{-2}$ and life time of 40 iterations, when new nodes are added. However, after the first addition, the added hidden nodes are selected from a large pool of candidate initialisations. The hidden nodes which correlate best with the remaining modelling error of the observations are selected. The first hidden nodes are able to model many of the large scale nonlinearities in the data. It is much more difficult to find useful hidden nodes by random initialisations later on since the new nodes tend to be quickly pruned away.

### 5.3 Post-nonlinear model

The evaluation of the cost function and the learning process of the PNFA model are similar to those of the general NFA. Approximating the nonlinearity is a little easier, as there is only one "source" input with larger posterior variance. The posterior variances of the MLP weights are typically smaller as a single weight affects and thus gains evidence from several observations. This allows for using a hybrid of a Gaussian quadrature with respect to the MLP inputs $y_{i}(t)=\sum_{j=1}^{M} A_{i j} s_{j}$ and a Taylor approximation with respect to the weights $\boldsymbol{\theta}_{f_{i}}$, as presented in Appendix A. The more general approach used with NFA could of course be used as well.

Using the approximation of the nonlinearity presented in Appendix A, the PNFA cost function can be evaluated using Eqs. (41)-(43) and (46)-(48). The
sources $\mathbf{s}$ and all the weights $\mathbf{A}$ and $\boldsymbol{\theta}_{f_{i}}$ are updated by conjugate gradient and all the other parameters by VB EM, as described for the NFA model in Sec. 5.1.2.

The initialisation of the PNFA method is similar to NFA: the sources are initialised with PCA. It would be possible to use other simple post-nonlinear methods such as Gaussianisation [27]. After the initialisation, the sources and the hyperparameters are kept fixed for the first 100 and 150 iterations, respectively, while the linear mapping $\mathbf{A}$ and the post-nonlinearities $f_{i}$ are updated.

### 5.4 Missing Values

The Bayesian framework facilitates principled handling of missing or partially missing elements in the data matrix [44]. Assuming the elements are missing at random, they provide no evidence and can hence be simply ignored in learning. After learning, the missing values can be reconstructed by using the posterior predictive distribution.

Missing values in nonlinear FA have been studied using NFA and HNFA models in [12,13].

## 6 Experiments

In this section, we present the results of two comparisons of the methods using artificial data sets. The artificial data sets are necessary for comparing different methods as the true sources are not known in most real applications and comparison of different methods is thus practically impossible.

### 6.1 General nonlinear mixtures

The NFA and HNFA methods were tested with an artificial example. The data set consisted of 1,000 samples from nonlinear mixtures of eight sources. Four of the sources were super-Gaussian and the remaining four were sub-Gaussian. The nonlinear mixing was a randomly initialised MLP network with sinh ${ }^{-1}$ as the activation function for the 30 hidden nodes. The standard deviation of the additive Gaussian noise was one tenth of that of the signal. This corresponds to a signal-to-noise ratio (SNR) of 20 dB . The same data set was also used in [10] and a noiseless version of it in [9].

Linear models (implemented as HNFA models without any hidden nodes) as
well as NFA and HNFA models were tested with varying number of sources. The values of the cost function attained after learning were compared. The best linear model had 14 sources while the best NFA and HNFA models had eight sources. Hence both NFA and HNFA methods are able to infer the correct subspace dimension, while the linear model tries to model nonlinear effects using extra dimensions.

In addition to the proposed variational methods, the experiment was repeated using MAP estimation for the NFA model (MAP-NFA). In this case all the priors and the model structure were fixed to the their generative values.

Unfortunately most other nonlinear BSS methods do not scale to problems of this size. For Kernel PCA [25] the size is not a problem, but out of several Gaussian and polynomial kernels tested, the best results were attained with a linear kernel which is equivalent to linear PCA. Finding a better kernel may certainly be possible, but at least very difficult. We also tried the MISEP method [22], but were unable to attain better results than using linear methods. Again, this may be caused by the authors' lack of experience with MISEP.


Fig. 3. Boxplot showing variation in the results of different algorithms in the general artificial mixture example.

The experiment was repeated a number of times with different random initialisations. Because of the flexibility of the nonlinear models, each of these ended in a different local optimum. The variation among the results is illustrated in Fig. 3 in terms of signal-to-noise ratio (SNR) with respect to the true sources after rotation of the extracted sources by symmetric FastICA [1]. The results shown are from 10 experiments with VB-NFA, 42 experiments with HNFA and 20 experiments with MAP-NFA. The results show that VB-NFA consistently finds very good solutions while there is lot more variation in the results of HNFA and especially MAP-NFA.

The figure also shows VB-NFA to be clearly superior to MAP-NFA, which
suffers from serious overfitting. Additionally, there is relatively strong correlation between the VB cost function value and SNR, making it easy to find the best solutions among the alternatives. For MAP-NFA, the correlation between posterior density value and SNR is significantly weaker.


Fig. 4. Each scatter plot shows the values of one original source signal plotted against the best corresponding estimated source signal after a rotation with FastICA. The perfect result would be a straight line.

After a nonlinear subspace has been estimated by NFA or HNFA, standard linear ICA algorithms [1] can be used for rotating the subspace to obtain independent source signals. Fig. 4 shows the scatter plots of the original sources and the sources obtained after a rotation by symmetric FastICA. The plots are from the best results in Fig. 3 according to the VB cost function, hence they are not necessarily the ones with the best overall SNR. On the first two rows, standard linear PCA has been used to estimate the subspace. The middle two rows show the results when the subspace was estimated using the HNFA method and the two bottom rows when the NFA method was used.

The results show that linear methods alone are not sufficient to separate the sources. The results attained by HNFA are reasonably good while those of NFA are already very good. The difference in performance between the methods may be partly due to the fact that the generative nonlinearity more closely resembles that used in NFA, even though it is not the same. The computation times required by NFA and HNFA in this experiment are comparable.

### 6.2 Post-nonlinear (PNL) mixtures

The proposed NFA and PNFA methods were also tested with a three-dimensional PNL mixture of two independent sources. The sources were a sine wave and uniformly distributed white noise. The PNL transformation used for generating the data contained two non-invertible post-nonlinear distortions:

$$
\mathbf{y}=\left[\begin{array}{cc}
1.2 & 0.2  \tag{54}\\
1 & 0.7 \\
0.2 & 0.8
\end{array}\right] \mathbf{s} ; \quad \mathbf{x}=\left[\begin{array}{c}
\left(y_{1}-0.5\right)^{2} \\
\left(y_{2}+0.4\right)^{2} \\
\tanh \left(2 y_{3}\right)
\end{array}\right]
$$

The observations were centred and normalised to unit variance and observation noise with variance 0.01 was added. The number of samples was 400 . The data set is illustrated in Fig. 5.

The NFA and PNFA models were trained by trying different model structures, i.e. different numbers of hidden nodes in the MLPs, and several random initialisations of weights of the linear and nonlinear mappings. The PNFA model with the lowest cost had 5 nodes in the hidden layers of all MLPs. The corresponding NFA model with the lowest cost had 10 hidden nodes in the single MLP. The PNFA model was trained for 10000 iterations and NFA for 5000 iterations.

The sources found by NFA and PNFA were further rotated by the FastICA algorithm to obtain independent signals (see Fig. 6). The scatter plots in Fig. 6a show how well the original sources were reconstructed. Each point


Fig. 5. The data set from the PNL mixture experiment.


Fig. 6. The sources found by the NFA (top) as well as PNFA (bottom) and further rotated with the FastICA algorithm. (a) The scatter plots; (b) the estimated time series; (c) the distribution of the sources. The signal-to-noise ratio is 10.92 dB for NFA and 12.95 dB for PNFA.
corresponds to one source $s_{i}(t)$. The abscissa of a point is the original source which was used for generating the data and the ordinate is the estimated source. The optimal result would be a straight line which would mean that the estimated values of the sources coincide with the true values. For PNFA, the sources were estimated quite well except for some points at the edges,


Fig. 7. The estimated post-nonlinear distortions $f_{i}$ against the functions used for generating the data (the dashed line). Each point in the figure corresponds to a single observation.
while the sources recovered by NFA are not as good.


Fig. 8. A grid of original sources $\mathbf{s}$ as mapped through the mixing to x and demixing to their reconstructions $\hat{\mathbf{s}}$ : the perfect solution would be a regular square. Because of the noise, the mappings are only approximate. The results are shown for the PNFA method (left) and Taleb-Jutten PNL ICA algorithm [19] based on non-invertible post-nonlinearities $f_{i}$ (right).

The PNL distortions learned by the best PNFA model are presented in Fig. 7. The post-nonlinearities $f_{i}$ are estimated quite well except for some points at the edges. The difficulties mostly affect the two quadratic functions which are difficult to model with such small MLP networks and relatively few observations, especially at those edges.

As expected, regular PNL ICA methods using non-invertible post-nonlinearities cannot recover the true sources at all. This is illustrated in Fig. 8, which shows a relatively clean separation by PNFA compared to severely tangled solution by Taleb-Jutten algorithm [19].

The mediocre overall quality of the results is understandable due to the great difficulty of the test problem: There are only two bounded sub-Gaussian sources in the mixture and their linear combinations are quite far from Gaussianity assumed by NFA and PNFA. Another difficulty is the complex PNL mapping with a small number of observations and several non-invertible post-
nonlinear distortions. Removing any of the observations from the mixture would make the mixing process non-injective and the separation problem unsolvable. The results of PNFA are slightly better than those of regular NFA, which is also natural as the data follows the more restricted model.

## 7 Discussion

Nonlinear BSS is a very difficult problem, as illustrated by the non-uniqueness results for nonlinear FA and ICA. From a Bayesian perspective, having a multitude of potential solutions that all should somehow be taken into account does not sound so unreasonable. The Bayesian averaging procedure together with reasonable priors on the mapping will help avoid most overly complicated solutions. This was also illustrated experimentally in Sec. 6.1. Using a slightly simpler EM-like approach with point estimates for the parameters of the nonlinear mapping would probably be enough to avoid most overfitting problems. However, as integration over the sources is the most complicated operation, the computational savings would be limited and the ability to infer the structure of the nonlinearity would be lost.

The variational approach used in this work has also other properties that further help regularise the problem. The flexible nonlinear model has many internal symmetries that are reflected in the true posterior. For purposes of the source separation problem, these symmetries are, however, not interesting. The variational approximation breaks the symmetry by ignoring the dependencies between different groups of parameters and the sources. The unimodal posterior approximation will find a broad region of potential solutions, thus returning essentially a single most plausible solution.

The regularisation power of the variational Bayesian approach is nicely demonstrated in [49], where HNFA is shown to prefer a linear model for a magnetoencephalographic (MEG) brain imaging data set where there are theoretical grounds to assume a linear mixing process. The example with approximately 50 sources extracted from 122 channels also demonstrates the scalability of HNFA, using NFA in a problem of this size would be infeasible.

Instead of a nonlinear model with a non-Gaussian prior capable of separating independent sources, simpler Gaussian source prior and post-processing by linear ICA has been used in this work to separate the independent sources. The main reason for using this suboptimal process is practical: even with a non-Gaussian source prior the factorial source posterior approximation would not allow determining the proper rotation of the independent sources [34]. This could be corrected by more advanced approximation techniques [5]. Application of such techniques for nonlinear models is an interesting topic of
future research.
In this work layered mappings have been used to model the nonlinearities. This gives a quite strong prior that allows us to solve the difficult example cases demonstrated in Sec. 6. The broad literature on applications of MLP networks on a wide variety of real world problems provides empirical evidence that such a prior may also be useful in real world applications. Still, having a more easily interpretable prior for the nonlinearities through a Gaussian process [26], for instance, would be interesting.

The presented HNFA and NFA methods have also been compared in reconstructing missing values in speech spectrograms in [33]. According to those experiments, the performance of HNFA was between linear FA and NFA when there were major differences. This seems to indicate that HNFA may emphasise less nonlinear solutions than NFA through stronger regularisation. Thus HNFA is likely better suited for mildly nonlinear problems.

## 8 Conclusion

We have presented a number of methods for BSS of nonlinear mixtures based on variational Bayesian learning. The basic NFA method uses a MLP network to model the nonlinearity. With a suitable linearisation technique, the variational Bayesian methodology can be easily applied to perform learning and inference on the model much more reliably than before [9]. The computational complexity of the resulting algorithm is quadratic in the number of the sources. To avoid this, an alternative HNFA method based on a hierarchical nonlinearity was proposed. The attained speedup comes at the expense of less accurate results as there are more latent variables in the model and the variational approximate posterior provides a poorer fit to the true posterior. The speeds of NFA and HNFA are comparable in the examples with approximately 10 sources presented in Sec. 6, but better scalability allows applying HNFA to larger problems, where no other alternatives are available. Additionally, a variant of NFA specifically designed for more restricted post-nonlinear mixtures was proposed.

Although HNFA seems less accurate in comparison to NFA in the experiment in Sec. 6.1, the same is not true in general. The comparison of Sec. 6.1 most likely favours NFA somewhat, because the structure of the generative nonlinearity is closer to that used in NFA, even though it is not the same. Other comparisons, such as some of the missing data experiments reported in [33] show HNFA outperform NFA.

Most of the proposed methods employ a Gaussian prior on the sources and are
thus unable to perform source separation by themselves. Still, post-processing the recovered sources by standard linear ICA was enough to recover the underlying independent sources in the presented examples. Using a non-Gaussian source model does not necessarily solve the problem, as ignoring the posterior correlations between the sources in the variational approximation may cause the model to prefer more orthogonal PCA-like solutions. Extending the model by using a Gaussian mixture source model and modelling the posterior correlations of the sources would be an interesting line of future work. The computational complexity of such a model is exponential in the number of the sources, but it would probably help in low dimensional cases.

In order to allow others use the presented methods more easily, free software implementations of the NFA and HNFA methods [39,41] are available on the web at http://www.cis.hut.fi/projects/bayes/software/.

## Acknowledgements

The authors wish to thank Tapani Raiko for useful comments and discussions and Tomas Östman for his help with the HNFA experiments. This work was supported in part by the IST Programme of the European Community, under the project BLISS, IST-1999-14190, and under the PASCAL Network of Excellence, IST-2002-506778. This publication only reflects the authors' views.

## A Approximating the MLP nonlinearity

The mean $\bar{f}_{i}$ and variance $\widetilde{f}_{i}$ of the outputs of the MLP used in Eq. (48) can be evaluated by multidimensional Gaussian integrals

$$
\begin{align*}
& \bar{f}_{i}(t)=\iint f_{i}\left(\mathbf{s}(t), \boldsymbol{\theta}_{\mathbf{f}}\right) q\left(\mathbf{s}(t), \boldsymbol{\theta}_{\mathbf{f}} \mid \boldsymbol{\xi}\right) d \mathbf{s}(t) d \boldsymbol{\theta}_{\mathbf{f}}  \tag{A.1}\\
& \tilde{f}_{i}(t)=\iint\left(f_{i}\left(\mathbf{s}(t), \boldsymbol{\theta}_{\mathbf{f}}\right)-\bar{f}_{i}(t)\right)^{2} q\left(\mathbf{s}(t), \boldsymbol{\theta}_{\mathbf{f}} \mid \boldsymbol{\xi}\right) d \mathbf{s}(t) d \boldsymbol{\theta}_{\mathbf{f}} . \tag{A.2}
\end{align*}
$$

These integrals depend on all the inputs and weights of the MLP network, thus leading to cases of the order of thousands of dimensions.

## A. 1 Multivariate nonlinearities

In [9], the moments were evaluated by using a first order Taylor approximation of $\mathbf{f}$ for (A.2) and second order approximation for (A.1). This method was
later found to produce unreliable estimates in cases of large source posterior variance, which are common when trying to extract a large number of sources. This caused instability of the algorithm in such cases. This can be corrected by a more accurate approximation based on global linearisation of $\mathbf{f}$. This can be implemented by replacing the derivatives of the activation function appearing in the Taylor scheme by global linearisations evaluated using Gauss-Hermite quadratures [15].

The Gauss-Hermite quadrature is a general method for approximating integrals of the form

$$
\begin{equation*}
I(\phi)=\int_{-\infty}^{\infty} \phi(y) N(y ; 0,1) d y \approx \sum_{i=1}^{n} w_{i} \phi\left(t_{i}\right), \tag{A.3}
\end{equation*}
$$

where $\phi: \mathbb{R} \rightarrow \mathbb{R}$ is a scalar function. For an approximation using $n$ points, the weights $w_{i}$ and abscissas $t_{i}$ can be selected so that the result is exact for all polynomials up to order $2 n$. We have used a 3 -point approximation as it provides a good compromise between accuracy and efficiency. The evaluation points can be easily transformed to handle general mean and variance of the input distribution to get general expectation of $\phi(y)$

$$
\begin{equation*}
\bar{\phi}(y)_{\mathrm{GH}}:=\sum_{i=1}^{n} w_{i} \phi\left(\bar{y}+t_{i} \sqrt{\widetilde{y}}\right) \approx\langle\phi(y)\rangle=\int_{-\infty}^{\infty} \phi(y) N(y ; \bar{y}, \widetilde{y}) d y . \tag{A.4}
\end{equation*}
$$

Variance of $\phi(y)$ can be evaluated through

$$
\begin{equation*}
\widetilde{\phi}(y)_{\mathrm{GH}}:=\sum_{i=1}^{n} w_{i}\left[\phi\left(\bar{y}+t_{i} \sqrt{\widetilde{y}}\right)-\bar{\phi}(y)_{\mathrm{GH}}\right]^{2} \approx\left\langle[\phi(y)-\langle\phi(y)\rangle]^{2}\right\rangle . \tag{A.5}
\end{equation*}
$$

Both the evaluated mean and variance of $\phi(y)$ depend on the mean and variance of $y$ in a nonlinear manner capable of taking into account the specific form of function $\phi$.

The quadrature depends on the mean and variance of $\mathbf{y}(t)=\mathbf{A s}(t)+\mathbf{a}$. These can be evaluated exactly because of the linearity of the mapping as

$$
\begin{align*}
\bar{y}_{i}(t) & =\sum_{j}\left(\bar{A}_{i j} \bar{s}_{j}(t)\right)+\bar{a}_{i}  \tag{A.6}\\
\widetilde{y}_{i, \text { tot }}(t) & =\sum_{j}\left(\widetilde{A}_{i j}\left(\bar{s}_{j}(t)^{2}+\widetilde{s}_{j}(t)\right)+\bar{A}_{i j}^{2} \widetilde{s}_{j}(t)\right)+\widetilde{a}_{i}, \tag{A.7}
\end{align*}
$$

where $\bar{\theta}$ denotes the mean and $\tilde{\theta}$ the variance of $\theta$. Here it is assumed that the posterior approximations $q\left(\boldsymbol{S} \mid \boldsymbol{\xi}_{S}\right)$ and $q\left(\boldsymbol{\theta}_{\mathbf{f}} \mid \boldsymbol{\xi}_{T T}\right)$ have diagonal covariances. Full covariances can be used instead without too much difficulty, if necessary. In order to get more accurate approximation, another quadrature is evaluated
using the variance of $y_{i}(t)$ originating mainly from $\boldsymbol{\theta}_{\mathbf{f}}$,

$$
\begin{equation*}
\widetilde{y}_{i, \text { weight }}(t)=\sum_{j} \widetilde{A}_{i j}\left(\bar{s}_{j}(t)^{2}+\widetilde{s}_{j}(t)\right)+\widetilde{a}_{i}, \tag{A.8}
\end{equation*}
$$

and using the implied $\widetilde{\phi}\left(y_{j}(t)\right)_{\text {GH, weight }}$ in the evaluation of the effects of these variances. The total variance (A.7) is still used in evaluation of the means and the evaluation of the effects of the variance of $\mathbf{s}(t)$.

The evaluated mean and variance can be used to define an effective linearisation of the hidden nodes by finding a corresponding linear function that would yield the same mean and variance. This yields the effective slope

$$
\begin{equation*}
\left\langle\phi^{\prime}\left(y_{i}(t)\right)\right\rangle:=\sqrt{\frac{\tilde{\phi}\left(y_{i}(t)\right)_{\mathrm{GH}}}{\widetilde{y}_{i}(t)}} . \tag{A.9}
\end{equation*}
$$

The effective linearisation is able to take into account the variance of the input, thus following the form of the function more globally when the variance is large.

Thus the final approximation for the mean is

$$
\begin{equation*}
\left.\bar{f}_{i}(t)\right\rangle=\sum_{j} \bar{B}_{i j} \bar{\phi}\left(y_{i}(t)\right)_{\mathrm{GH}}+\bar{b}_{i} \tag{A.10}
\end{equation*}
$$

and for the variance

$$
\begin{array}{r}
\left\langle\left(f_{i}(t)-\bar{f}_{i}(t)\right)^{2}\right\rangle=\sum_{j}\left(\widetilde{B}_{i j}\left(\bar{\phi}(y)_{\mathrm{GH}}^{2}+\widetilde{\phi}\left(y_{j}(t)\right)_{\mathrm{GH}}\right)+\bar{B}_{i j}^{2} \widetilde{\phi}\left(y_{j}(t)\right)_{\mathrm{GH}, \text { weight }}\right) \\
+\widetilde{b}_{i}+\widetilde{\nabla}_{\mathbf{s}} f_{i} \operatorname{diag}(\widetilde{\mathbf{s}}) \widetilde{\nabla}_{\mathbf{s}} f_{i}^{T} \tag{A.11}
\end{array}
$$

where the effective gradient is

$$
\begin{equation*}
\widetilde{\nabla}_{s_{j}} f_{i}(t)=\sum_{l} \bar{B}_{i l}\left\langle\phi^{\prime}\left(y_{l}(t)\right)\right\rangle \bar{A}_{l j} . \tag{A.12}
\end{equation*}
$$

## A. 2 Scalar post-nonlinearities

In case of the MLP networks modelling the scalar post-nonlinearities in PNFA, the approximation is slightly easier as it is possible to directly use the GaussHermite quadrature with respect to the inputs of the MLP. As the variances of the weights are usually small, their effects are represented sufficiently well by using first-order Taylor approximation of the network with respect to them. Thus the mean of the output is approximated as

$$
\begin{equation*}
\bar{f}_{i}(t)=\sum_{j} w_{j} f_{i}\left(\hat{y}_{j}(t), \overline{\boldsymbol{\theta}}_{\mathbf{f}_{i}}\right) \tag{A.13}
\end{equation*}
$$

where $\mathbf{y}=\mathbf{A s}, w_{j}$ are the weights and $\hat{y}_{j}(t)=\bar{y}(t)+t_{j} \sqrt{y(t)}$ are the basis points of the Gauss-Hermite quadrature corresponding to the abscissas $t_{j}$, and $\overline{\boldsymbol{\theta}}_{\mathbf{f}_{i}}$ denotes the mean of the weights $\boldsymbol{\theta}_{\mathbf{f}_{i}}$.

Correspondingly, the variance is approximated by a combined Gauss-Hermite and Taylor approximation

$$
\begin{align*}
\left\langle\left(f_{i}(t)-\bar{f}_{i}(t)\right)^{2}\right\rangle= & \sum_{j} w_{j}\left[\left(f_{i}\left(\hat{y}_{j}(t), \overline{\boldsymbol{\theta}}_{\mathbf{f}_{i}}\right)-\bar{f}_{i}(t)\right)^{2}\right. \\
& \left.+\nabla_{\boldsymbol{\theta}_{\mathbf{f}_{i}}} f_{i}\left(\hat{y}_{j}(t), \overline{\boldsymbol{\theta}}_{\mathbf{f}_{i}}\right) \operatorname{diag}\left(\widetilde{\boldsymbol{\theta}}_{\mathbf{f}_{i}}\right) \nabla_{\boldsymbol{\theta}_{\mathbf{f}_{i}}} f_{i}\left(\hat{y}_{j}(t), \overline{\boldsymbol{\theta}}_{\mathbf{f}_{i}}\right)^{T}\right] . \tag{A.14}
\end{align*}
$$

## References

[1] A. Hyvärinen, J. Karhunen, E. Oja, Independent Component Analysis, J. Wiley, 2001.
[2] H. Attias, Independent factor analysis, Neural Computation 11 (4) (1999) 803851.
[3] H. Lappalainen, Ensemble learning for independent component analysis, in: Proc. Int. Workshop on Independent Component Analysis and Signal Separation (ICA'99), Aussois, France, 1999, pp. 7-12.
[4] J. Miskin, Ensemble learning for independent component analysis, Ph.D. thesis, University of Cambridge, UK (2000).
[5] P. Højen-Sørensen, O. Winther, L. K. Hansen, Mean-field approaches to independent component analysis, Neural Computation 14 (4) (2002) 889-918.
[6] R. A. Choudrey, S. J. Roberts, Variational mixture of Bayesian independent component analyzers, Neural Computation 15 (1) (2003) 213-252.
[7] K. Chan, T.-W. Lee, T. J. Sejnowski, Variational Bayesian learning of ICA with missing data, Neural Computation 15 (8) (2003) 1991-2011.
[8] H. Attias, Independent factor analysis with temporally structured sources, in: S. Solla, T. Leen, K.-R. Müller (Eds.), Advances in Neural Information Processing Systems 12, MIT Press, Cambridge, MA, USA, 2000, pp. 386-392.
[9] H. Lappalainen, A. Honkela, Bayesian nonlinear independent component analysis by multi-layer perceptrons, in: M. Girolami (Ed.), Advances in Independent Component Analysis, Springer-Verlag, Berlin, 2000, pp. 93-121.
[10] H. Valpola, T. Östman, J. Karhunen, Nonlinear independent factor analysis by hierarchical models, in: Proc. 4th Int. Symp. on Independent Component Analysis and Blind Signal Separation (ICA2003), Nara, Japan, 2003, pp. 257262.
[11] A. Ilin, A. Honkela, Postnonlinear independent component analysis by variational Bayesian learning, in: C. G. Puntonet, A. Prieto (Eds.), Proc. of the Fifth Int. Conf. on Independent Component Analysis and Blind Signal Separation (ICA 2004), Vol. 3195 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Granada, Spain, 2004, pp. 766-773.
[12] T. Raiko, H. Valpola, Missing values in nonlinear factor analysis, in: Proc. of the 8th Int. Conf. on Neural Information Processing (ICONIP'01), Shanghai, 2001, pp. 822-827.
[13] T. Raiko, H. Valpola, T. Östman, J. Karhunen, Missing values in hierarchical nonlinear factor analysis, in: Proc. of the Int. Conf. on Artificial Neural Networks and Neural Information Processing (ICANN/ICONIP 2003), Istanbul, Turkey, 2003, pp. 185-189.
[14] A. Honkela, S. Harmeling, L. Lundqvist, H. Valpola, Using kernel PCA for initialisation of variational Bayesian nonlinear blind source separation method, in: C. G. Puntonet, A. Prieto (Eds.), Proc. of the 5th Int. Conf. on Independent Component Analysis and Blind Signal Separation (ICA 2004), Vol. 3195 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Granada, Spain, 2004, pp. 790-797.
[15] A. Honkela, H. Valpola, Unsupervised variational Bayesian learning of nonlinear models, in: L. Saul, Y. Weiss, L. Bottou (Eds.), Advances in Neural Information Processing Systems 17, MIT Press, Cambridge, MA, USA, 2005, pp. 593-600.
[16] C. Jutten, J. Karhunen, Advances in blind source separation (BSS) and independent component analysis (ICA) for nonlinear mixtures, International Journal of Neural Systems 14 (5) (2004) 267-292.
[17] L. B. Almeida, Nonlinear Source Separation, Synthesis Lectures on Signal Processing, Morgan\&Claypool Publishers, 2006.
[18] A. Hyvärinen, P. Pajunen, Nonlinear independent component analysis: Existence and uniqueness results, Neural Networks 12 (3) (1999) 429-439.
[19] A. Taleb, C. Jutten, Source separation in post-nonlinear mixtures, IEEE Trans. on Signal Processing 47 (10) (1999) 2807-2820.
[20] S. Haykin, Neural Networks - A Comprehensive Foundation, 2nd ed., PrenticeHall, 1999.
[21] S. Hochreiter, J. Schmidhuber, LOCOCODE performs nonlinear ICA without knowing the number of sources, in: Proc. Int. Workshop on Independent Component Analysis and Blind Signal Separation (ICA'99), Aussois, France, 1999, pp. 149-154.
[22] L. B. Almeida, MISEP - linear and nonlinear ICA based on mutual information, Journal of Machine Learning Research 4 (2003) 1297-1318.
[23] L. B. Almeida, Separating a real-life nonlinear image mixture, Journal of Machine Learning Research 6 (2005) 1199-1229.
[24] M. S. C. Almeida, H. Valpola, J. Särelä, Separation of nonlinear image mixtures by denoising source separation, in: Proc. 6th Int. Conf. on Independent Component Analysis and Blind Source Separation (ICA 2006), Charleston, South Carolina, USA, 2006, pp. 8-15.
[25] B. Schölkopf, A. Smola, K.-R. Müller, Nonlinear component analysis as a kernel eigenvalue problem, Neural Computation 10 (5) (1998) 1299-1319.
[26] N. Lawrence, Probabilistic non-linear principal component analysis with Gaussian process latent variable models, Journal of Machine Learning Research 6 (2005) 1783-1816.
[27] A. Ziehe, M. Kawanabe, S. Harmeling, K.-R. Müller, Blind separation of post-nonlinear mixtures using linearizing transformations and temporal decorrelation, Journal of Machine Learning Research 4 (2003) 1319-1338.
[28] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, Neural Networks 2 (5) (1989) 359-366.
[29] H. Valpola, Nonlinear independent component analysis using ensemble learning: theory, in: Proc. Int. Workshop on Independent Component Analysis and Blind Signal Separation (ICA2000), Helsinki, Finland, 2000, pp. 251-256.
[30] M. Stinchcombe, H. White, Universal approximation using feedforward networks with non-sigmoid hidden layer activation functions, in: Proc. International Joint Conference on Neural Networks (IJCNN '89), 1989, pp. I-613-617.
[31] H. Valpola, T. Raiko, J. Karhunen, Building blocks for hierarchical latent variable models, in: Proc. 3rd Int. Conf. on Independent Component Analysis and Signal Separation (ICA2001), San Diego, USA, 2001, pp. 710-715.
[32] M. Harva, T. Raiko, A. Honkela, H. Valpola, J. Karhunen, Bayes Blocks: An implementation of the variational Bayesian building blocks framework, in: Proc. of the 21st Conf. on Uncertainty in Artificial Intelligence (UAI 2005), Edinburgh, Scotland, 2005, pp. 259-266.
[33] T. Raiko, H. Valpola, M. Harva, J. Karhunen, Building blocks for variational Bayesian learning of latent variable models, Tech. Rep. E4, Helsinki University of Technology, Publications in Computer and Information Science, Espoo, Finland, available at http://www.cis.hut.fi/Publications/. Accepted for publication in Journal of Machine Learning Research conditioned on minor revisions. (2006).
[34] A. Ilin, H. Valpola, On the effect of the form of the posterior approximation in variational learning of ICA models, Neural Processing Letters 22 (2) (2005) 183-204.
[35] S. Ikeda, ICA on noisy data: A factor analysis approach, in: M. Girolami (Ed.), Advances in Independent Component Analysis, Springer-Verlag, Berlin, 2000, pp. 201-215.
[36] T. S. Jaakkola, Tutorial on variational approximation methods, in: M. Opper, D. Saad (Eds.), Advanced Mean Field Methods: Theory and Practice, The MIT Press, Cambridge, MA, USA, 2001, pp. 129-159.
[37] J. Winn, C. M. Bishop, Variational message passing, Journal of Machine Learning Research 6 (2005) 661-694.
[38] H. Lappalainen, J. Miskin, Ensemble learning, in: M. Girolami (Ed.), Advances in Independent Component Analysis, Springer-Verlag, Berlin, 2000, pp. 75-92.
[39] A. Honkela, H. Valpola, X. Giannakopoulos, Nonlinear factor analysis Matlab package, http://www.cis.hut.fi/projects/bayes/software/ (2004).
[40] K. Fukumizu, S. Amari, Local minima and plateaus in hierarchical structures of multilayer perceptrons, Neural Networks 13 (3) (2000) 317-327.
[41] H. Valpola, A. Honkela, M. Harva, A. Ilin, T. Raiko, T. Östman, Bayes Blocks software library, http://www.cis.hut.fi/projects/bayes/ software/ (2003).
[42] A. Honkela, H. Valpola, J. Karhunen, Accelerating cyclic update algorithms for parameter estimation by pattern searches, Neural Processing Letters 17 (2) (2003) 191-203.
[43] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann, 1988.
[44] T. Raiko, Partially observed values, in: Proc. Int. Joint Conf. on Neural Networks (IJCNN’04), Budapest, Hungary, 2004, pp. 2825-2830.
[45] A. Honkela, H. Valpola, Variational learning and bits-back coding: an information-theoretic view to Bayesian learning, IEEE Transactions on Neural Networks 15 (4) (2004) 800-810.
[46] H. Valpola, J. Karhunen, An unsupervised ensemble learning method for nonlinear dynamic state-space models, Neural Computation 14 (11) (2002) 2647-2692.
[47] T. Kohonen, Self-Organizing Maps, 3rd Edition, Springer, 2001.
[48] J. Vesanto, J. Himberg, E. Alhoniemi, J. Parhankangas, SOM toolbox for Matlab 5, Tech. Rep. A57, Helsinki University of Technology, Publications in Computer and Information Science, Espoo, Finland, URL: http://www.cis. hut.fi/projects/somtoolbox/ (2000).
[49] A. Honkela, T. Östman, R. Vigário, Empirical evidence of the linear nature of magnetoencephalograms, in: Proc. 13th European Symposium on Artificial Neural Networks (ESANN 2005), Bruges, Belgium, 2005, pp. 285-290.


[^0]:    * Corresponding author. Tel. +358-9-451 5283, Fax +358-9-451 3277

    Email address: antti.honkela@tkk.fi (Antti Honkela).
    URL: http://www.cis.hut.fi/projects/bayes/ (Antti Honkela).

[^1]:    ${ }^{1}$ At the time his surname was Lappalainen.

[^2]:    $\overline{3}$ A Matlab toolbox [39] implementing the NFA method is available at http:// www.cis.hut.fi/projects/bayes/software/.

